Chapter 11

IS-IS Configuration Guidelines

To configure Intermediate System to Intermediate System (IS-IS), you include statements at the [edit protocols isis] hierarchy level of the configuration. For routing instances, include the statements at the [edit routing-instances *routing-instance-name* protocols isis]. You can include the following statements in the configuration:

```
protocols {
    isis {
        disable;
        authentication-key key;
        authentication-type authentication;
        export [ policy-names ];
        ignore-attached-bit;
        graceful-restart {
            disable;
        }
        label-switched-path name level level metric metric;
        level level-number {
            authentication-key key;
            authentication-type authentication;
            external-preference preference;
            no-csnp-authentication;
            no-hello-authentication;
            no-psnp-authentication;
            preference preference;
            wide-metrics-only;
        }
        lsp-lifetime seconds;
        multicast-topology;
        no-authentication-check;
        overload <timeout seconds>;
        reference-bandwidth reference-bandwidth;
        rib-group group-name;
        spf-delay milliseconds;
        traffic-engineering {
            disable;
            shortcuts;
        }
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
```

```
        interface interface-name {
            authentication-key key;
            authentication-type authentication;
            disable;
            checksum;
            csnp-interval (seconds | disable);
            hello-authentication-key key;
            hello-authentication-type authentication;
            lsp-interval milliseconds;
            mesh-group (value | blocked);
            passive;
            level level-number {
                authentication-key key;
                authentication-type authentication;
                disable;
                hello-authentication-key key;
                hello-authentication-type authentication;
                hello-interval seconds;
                hold-time seconds;
                metric metric;
                passive;
                priority number;
                te-metric metric;
            }
        }
    }
}
```

By default, IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which an ISO address is configured.

This chapter discusses the following topics that provide information about configuring IS-IS:

See page 173 for an IS-IS configuration example.

Minimum IS-IS Configuration

For IS-IS to run on the router, you must enable IS-IS on the router, configure a network entity title (NET) on one of the router's interfaces (preferably the loopback interface, lo0), and configure the ISO family on all interfaces on which you want IS-IS to run. When you enable IS-IS, Level 1 and Level 2 are enabled by default. The following is the minimum IS-IS configuration. In the address statement, *address* is the NET.

```
interfaces {
    lo0 {
        unit logical-unit-number {
            family iso {
                address address;
            }
        }
    }
    type-fpc/pic/port {
        unit logical-unit-number {
            family iso;
        }
    }
}
protocols {
    isis {
        interface all;
    }
}
```

When you configure IS-IS on an interface, you must also include the family iso statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. For more information about the family iso statement, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Configure IS-IS Authentication

All IS-IS protocol exchanges can be authenticated to guarantee that only trusted routers participate in the autonomous system (AS) routing. By default, IS-IS authentication is disabled on the router.

To configure IS-IS authentication, you must define an authentication password and specify the authentication type.

You can configure one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. Simple authentication is included for compatibility with existing IS-IS implementations. However, we recommend that you do *not* use this authentication method because it is insecure (the text can be "sniffed").

- HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving router uses an authentication key (password) to verify the packet. HMAC-MD5 authentication is defined in Requests for Comments (RFC) 2104. Note that this RFC presents only a proposal for using HMAC-MD5 with IS-IS; it is currently not a standard.

You can also configure more fine-grained authentication for hello packets. To do this, see "Configure Authentication for Hello Packets" on page 183.

To enable authentication and specify an authentication method, include the authentication-type statement, specifying the simple or md5 authentication type:

```
authentication-type authentication;
```

You can enable authentication for all IS-IS levels (at the [edit protocols isis] hierarchy level), for an individual level (at the [edit protocols isis level *level-number*] hierarchy level), for an individual interface (at the [edit protocols isis interface *interface-name*] hierarchy level), and for an individual level on an interface (at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level. For routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level, the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level, the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level, and the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level.

To configure a password, include the authentication-key statement. The authentication password for all routers in a domain must be the same.

```
authentication-key key;
```

You can configure a password for all IS-IS levels (at the [edit protocols isis] hierarchy level), for an individual level (at the [edit protocols isis level *level-number*] hierarchy level), for an individual interface (at the [edit protocols isis interface *interface-name*] hierarchy level), and for an individual level on an interface (at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level. For routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level, the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level, the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level, and the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level.

The password can contain up to 255 characters. If you include spaces, enclose all characters in quotation marks (" ").

If you are using the JUNOS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces that are shared with a JUNOS implementation.

Authentication of hello packets, partial sequence number PDU (PSNP), and complete sequence number PDU (CSNP) may be suppressed to enable interoperability with the routing software of different vendors. Different vendors handle authentication in various ways, and suppressing authentication for different PDU types may be the simplest way to allow compatibility within the same network.

To configure IS-IS to generate authenticated packets, but not to check the authentication on received packets, include the no-authentication-check statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
no-authentication-check;
```

To suppress authentication of IS-IS hello packets, include the no-hello-authentication statement at the [edit protocols isis level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level):

```
[edit protocols isis]
no-hello-authentication;
```

To suppress authentication of PSNP packets, include the no-psnp-authentication statement at the [edit protocols isis level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level):

```
[edit protocols isis]
no-psnp-authentication;
```

To suppress authentication of CSNP packets, include the no-csnp-authentication statement at the [edit protocols isis level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level):

```
[edit protocols isis]
no-csnp-authentication;
```

The authentication and the no-authentication statements must to be configured at the same hierarchy level. Configuring authentication at the interface hierarchy level and configuring no-authentication at the isis hierarchy level has no effect.

Example: Migrating a Domain to Use Authentication

Having IS-IS not check authentication is useful when you have a domain that does not use authentication and are migrating it so that it uses authentication. To ensure a smooth migration, first configure all routers so that they generate authenticated packets but do not check the authentication of received packets (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols]
isis {
    authentication-key key;
    authentication-type authentication;
    no-authentication-check;
}
```

Check the network to ensure that no problems are occurring as a result of the generated authenticated packets. Check for interoperability problems, and also check log messages to identify routers that are using a password that differs from the domain-wide password.

When you are sure that no problems are occurring, reconfigure the routers, removing the no-authentication-check statement:

```
[edit protocols]
isis {
    authentication-key key;
    authentication-type authentication;
}
```

Configure Interface-Specific Properties

You can configure interface-specific IS-IS properties by including the interface statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level). These properties are explained later in this chapter.

```
[edit protocols isis]
interface interface-name {
    authentication-key key;
    authentication-type authentication;
    disable;
    checksum;
    csnp-interval (seconds | disable);
    hello-authentication-key key;
    hello-authentication-type authentication;
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    passive;
```

```
    level level-number {
        authentication-key key;
        authentication-type authentication;
        disable;
        hello-authentication-type authentication;
        hello-authentication-key key;
        hello-interval seconds;
        hold-time seconds;
        metric metric;
        passive;
        priority number;
        te-metric metric;
    }
}
```

For *interface-name*, specify the full interface name, including the physical and logical address components. To configure all interfaces, specify the interface name as all. For details about configuring interfaces, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Enable Checksum

You can enable checksum for packets on a per-interface basis. To enable checksum, include the checksum statement at the [edit protocols isis interface *interface-name*] hierarchy level:

```
[edit protocols isis interface interface-name]
checksum;
```

Configure the CSNP Interval

By default, IS-IS sends complete sequence number (CSN) packets periodically. If the router is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the router is on a point-to-point interface, it sends CSN packets every 5 seconds. You might want to modify the default interval to protect against link-state PDU (LSP) flooding.

To modify the CSNP interval, include the csnp-interval statement at the [edit protocols isis interface *interface-name*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
csnp-interval seconds;
```

The time can range from 1 through 65,535 seconds.

To configure the interface not to send any CSN packets, specify the disable option:

```
[edit protocols isis interface interface-name]
csnp-interval disable;
```

Configure Mesh Groups

A *mesh group* is a set of routers that are fully connected; that is, they have a fully meshed topology. When LSP packets are being flooded throughout an area, each router within a mesh group receives only a single copy of an LSP packet instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of LSP packets.

To create a mesh group and designate that an interface is part of the group, assign a mesh-group number to all the router interfaces in the group (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
mesh-group value;
```

To prevent an interface in the mesh group from flooding LSPs, configure blocking on that interface:

```
[edit protocols isis interface interface-name]
mesh-group blocked;
```

Modify the Interface Metric

All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.

The cost of a route is described by a single dimensionless metric that is determined using the following formula:

```
cost = reference-bandwidth / bandwidth
```

reference-bandwidth is the reference bandwidth. If the reference bandwidth is not configured, all interfaces have a default metric of 10 (with the exception of the lo0 interface, which has a default metric of 0).

To modify the reference bandwidth, include the reference-bandwidth statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
reference-bandwidth reference-bandwidth;
```

For example, if you set the reference bandwidth to 1 Gbps (that is, *reference-bandwidth* is set to 1,000,000,000), a 100-Mbps interface has a default metric of 10.

For additional information about IS-IS interface metrics, see “Modify the IS-IS Metric” on page 185.

Enable Wide Metrics for Traffic Engineering

Normally, IS-IS metrics can have values up to 63, and IS-IS generates two Type Length Values (TLVs), one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 2^{24}-1 (16,777,215).

By default, the JUNOS software allows a maximum metric value of 63 and generates both pairs of TLVs. To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the wide-metrics-only statement at the [edit protocols isis level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis level]
wide-metrics-only;
```

By default, the JUNOS software supports the sending and receiving of wide metrics.

Configure Route Preferences

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected. For more information about route preferences, see “Route Preferences” on page 6.

By default, Level 1 IS-IS internal routes have a preference value of 15, Level 2 IS-IS internal routes have a preference of 18, Level 1 IS-IS external routes have a preference of 160, and Level 2 external routes have a preference of 165. To change the preference values, include the preference statement (for internal routes) or the external-preference statement (for external routes) at the [edit protocols isis level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis level *level-number*] hierarchy level):

```
[edit protocols isis level level-number]
external-preference preference;
preference preference;
```

The preference value can range from 0 through 255.

Configure IS-IS Levels on an Interface

You can administratively divide a single AS into smaller groups called *areas*. There are two types of areas: Level 1 areas and Level 2 areas. Routers in Level 1 areas route within the area and, when the destination is outside the area, toward a Level 2 router. Routers in Level 2 areas route between areas and toward other ASs.

You configure each router interface to be in an area. Any interface can be in any area.

You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.

To configure an area, include the level statement at the [edit protocols isis interface *interface-name*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
level level-number {
    authentication-key key;
    authentication-type authentication;
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    metric metric;
    passive;
    priority number;
    te-metric metric;
}
```

The statements within the level statement allow you to configure the following optional level-specific properties:

Disable IS-IS on a Level on page 182

Advertise Interface Addresses without Running IS-IS on page 183

Configure Authentication for Hello Packets on page 183

Modify the Hello Interval on page 184

Modify the Hold-Time Value on page 185

Modify the IS-IS Metric on page 185

Modify the Traffic Engineering Metric on page 185

Configure the Priority for Becoming the Designated Router on page 186

Configure the Router to Advertise without Running IS-IS on page 186

Disable IS-IS on a Level

By default, IS-IS is enabled for Level 1 and Level 2 areas on all enabled interfaces on which the iso protocol family is enabled (at the [edit interfaces *interface* unit *logical-unit-number*] hierarchy level). To disable IS-IS at any particular level on an interface, include the disable statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
disable;
```

Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.

Example: Disable IS-IS on a Level

On SONET/SDH interface so-0/0/0, enable IS-IS for Level 1 only. With this configuration, tracing messages periodically will indicate that IS-IS is creating Level 2 LSPs. However, because IS-IS for Level 2 is disabled, these LSPs are never distributed to neighboring routers.

```
protocols {
    isis {
        traceoptions {
            file isis size 1m files 10;
            flag spf;
            flag lsp;
            flag error;
        }
        interface so-0/0/0 {
            level 2 {
                disable;
            }
        }
    }
}
```

Advertise Interface Addresses without Running IS-IS

By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level. To advertise the direct interface addresses without actually running IS-IS on that interface or level, include the passive statement at the [edit protocols isis interface *interface-name*] or [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] or [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
passive;
```

Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.

Configure Authentication for Hello Packets

You can configure authentication for all IS-IS hello packets for an interface and, to achieve a more-fine grained authentication, you can configure authentication for a given IS-IS level on that interface. If you configure a point-to-point link and if you enable both levels, the hello packets are sent with the password configured for Level 1.

By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

To configure IS-IS hello packet authentication, you must define an authentication password and specify the authentication type.

To enable hello authentication for an interface, include the hello-authentication statement at the [edit protocols isis interface *interface-name*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
hello-authentication-type authentication;
```

To configure the password, include the hello-authentication-key statement at the [edit protocols isis interface *interface-name*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
hello-authentication-key key;
```

To enable hello authentication at an IS-IS level on an interface, include the hello-authentication statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
hello-authentication-type authentication;
```

To configure a password at an IS-IS level on an interface, include the hello-authentication-key statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
hello-authentication-key key;
```

Modify the Hello Interval

Routers send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet. By default, a designated intersystem (DIS) router sends hello packets every 3 seconds, and a non-DIS router sends hello packets every 9 seconds.

To modify how often the router sends hello packets out of an interface, include the hello-interval statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
hello-interval seconds;
```

Modify the Hold-Time Value

The hold time specifies how long a neighbor should consider this router to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this router within the hold time, it marks the router as being unavailable. The default hold-time value is three times the default hello interval: 9 seconds for a DIS router and 27 seconds for a non-DIS router.

To modify the hold-time value on the local router, include the hold-time statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
hold-time seconds;
```

Modify the IS-IS Metric

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 2^{24}-1 (16,777,215) if you are using wide metrics. The default metric value is 10 (with the exception of the lo0 interface, which has a default metric of 0). To modify the default value, include the metric statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
metric metric;
```

For more information about IS-IS interface metrics, see “Modify the Interface Metric” on page 180.

Modify the Traffic Engineering Metric

When traffic engineering is enabled on the router, you can configure an IS-IS metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the traffic engineering database (TED). Its value does not affect normal IS-IS forwarding.

To modify the default value, include the te-metric statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
te-metric metric;
```

Configure the Priority for Becoming the Designated Router

A router advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This router is responsible for sending network link-state advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area.

The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127; routers with a higher value are more likely to become the designated router. By default, routers have a priority value of 64.

To modify the interface's priority value, include the priority statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*] hierarchy level):

```
[edit protocols isis interface interface-name level level-number]
priority number;
```

Configure the Router to Advertise without Running IS-IS

The router can advertise the direct interface addresses on an interface or on a sub-level of the interface without actually running IS-IS on that interface or at that level. This occurs in passive mode.

To enable an interface as passive, include the passive statement at the [edit protocols isis interface *interface-name* level *level-number*] hierarchy level:

```
[edit protocols isis interface interface-name level level-number]
passive;
```

Modify the LSP Interval

By default, the router sends one LSP packet out an interface every 100 milliseconds. To modify this interval, include the lsp-interval statement at the [edit protocols isis interface *interface-name*] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis interface *interface-name*] hierarchy level):

```
[edit protocols isis interface interface-name]
lsp-interval milliseconds;
```

To disable the transmission of all LSP packets, set the interval to 0.

Modify the LSP Lifetime

By default, link-state PDUs (LSPs) are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the LSP lifetime, normally is sufficient to guarantee that LSPs never expire.

To modify the LSP lifetime, include the lsp-lifetime statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
lsp-lifetime seconds;
```

The time can range from 350 to 65,535 seconds.

The LSP refresh interval is derived from the LSP lifetime and is equal to the lifetime minus 317 seconds.

Advertise Label-Switched Paths into IS-IS

You can advertise label-switched paths into IS-IS as point-to-point links, and the label-switched paths can be used in SPF calculations. The advertisement contains a local address (the from address of the label-switched path), a remote address (the to address of the label-switched path), and a metric with the following precedence:

Use the label-switched path metric defined under IS-IS.

Use the label-switched path metric configured for the label-switched path under MPLS.

If you do not configure any of the above, use the default IS-IS metric of 10.

To advertise label-switched paths, include the label-switched-path statement, with a specified level and metric, at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
label-switched-path name level level metric metric;
```

Before a single-hop label-switched path between a multi-access link can be announced as up and used in SPF calculations, you must configure a label-switched path in both directions between two label-switched routers.

For more information about advertising label-switched paths, see the *JUNOS Software Configuration Guide: MPLS Applications*.

Configure the Router to Appear Overloaded

You can configure the local router so that it appears to be overloaded. You might want to do this when you want the router to participate in IS-IS routing, but do not want it to be used for transit traffic. (Note that traffic to immediately attached interfaces continues to transit the router.) To mark the router as overloaded, include the overload statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
overload;
```

To specify the number of seconds at which overload is reset, include the timeout option when specifying the overload statement:

```
[edit protocols isis]
overload timeout <seconds>;
```

The time can range from 60 through 1,800 seconds.

Configure the SPF Delay

You can configure the shortest-path-first (SPF) algorithm delay. The SPF algorithm delay is the amount of time between the detection of a topology change and when the SPF algorithm actually runs to achieve convergence. The shorter the delay, the shorter the convergence time.

To configure the SPF delay, include the spf-delay statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
spf-delay milliseconds;
```

The time can range from 50 through 1000 milliseconds.

Disable Graceful Restart

Graceful restart allows a router to restart with minimal effects to the network, and is enabled globally for all routing protocols at the [edit routing-options] hierarchy level. When graceful restart for IS-IS is enabled, the restarting router is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.

To disable graceful restart for IS-IS, include the disable statement at the [edit protocols isis graceful-restart] hierarchy level:

```
[edit protocols isis]
graceful-restart {
    disable;
}
```

IS-IS and Multipoint Configurations

IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.

Configure IS-IS Traffic Engineering Attributes

You can configure two IS-IS traffic engineering attributes:

Configure IS-IS to Use IGP Shortcuts on page 189

Disable IS-IS Support for Traffic Engineering on page 190

When configuring traffic engineering support, you can also configure IS-IS to use metric values greater than 63, as described in “Enable Wide Metrics for Traffic Engineering” on page 181.

Configure IS-IS to Use IGP Shortcuts

IS-IS always performs SPF calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS also installs the prefix in the inet.3 routing table. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path’s (LSP’s) egress router. BGP uses this routing table to resolve next-hop addresses.

If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the label-switched path as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses a label-switched path along the path to reach the egress router.

To configure IS-IS so that it uses label-switched paths as shortcuts when installing information in the inet.3 routing table, include the shortcuts statement at the [edit protocols isis traffic-engineering] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis traffic-engineering] hierarchy level):

```
[edit protocols isis]
traffic-engineering {
    shortcuts;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure label-switched path shortcuts only on these routers.

For more information about configuring label-switched paths and MPLS, see the *JUNOS Internet Softw are Configur ation Guide: MPLS Applications* .

Disable IS-IS Support for Traffic Engineering

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database (TED). To disable this support, and to disable IS-IS shortcuts if they are configured, include the disable statement at the [edit protocols isis traffic-engineering] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis traffic-engineering] hierarchy level):

```
[edit protocols isis]
traffic-engineering {
    disable;
}
```

Disable IS-IS on the Router

To disable IS-IS on the router without removing the IS-IS configuration statements from the configuration, include the disable statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols]
isis {
    disable;
}
```

To re-enable IS-IS, remove the disable statement from the configuration:

```
[edit protocols]
user@host# delete isis disable
[edit protocols]
user@host# show
isis;
```

Configure IS-IS Routing Policy

All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.

For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a *routing policy* for that protocol. For information about defining routing policy, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

For IS-IS, you can apply routing policies that affect how routing protocol process (rpd) exports routes into IS-IS. To do this, include the export statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
export [ policy-names ];
```

For IS-IS, you should not apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol could easily lead to an inconsistent topology database.

Examples: Configure IS-IS Routing Policy

Define a policy that allows only host routes from USC (128.125.0.0/16), and apply the policy to routes exported from the routing table into IS-IS:

```
policy-options {
    policy-statement usc-hosts-only {
        term first {
            from {
                route-filter 128.125.0.0/16 upto /31;
            }
            then reject;
        }
        then accept;
    }
}
protocols {
    isis {
        export usc-hosts-only;
    }
}
```

Define a policy that takes Border Gateway Protocol (BGP) routes from the Edu community and places them into IS-IS with a metric of 14. Apply the policy to routes exported from the routing table into IS-IS:

```
protocols {
    isis {
        export edu-to-isis;
    }
}
policy-options {
    community Edu members 666:5;
    policy-statement edu-to-isis {
        from {
            protocol bgp;
            community Edu;
        }
        to protocol isis;
        then metric 14;
    }
}
```

Define a policy that rejects all IS-IS Level 1 routes so that none are exported into IS-IS:

```
policy-options {
    policy-statement level1 {
        term first {
            from level 1;
            then reject;
        }
        then accept;
    }
}
protocols {
    isis {
        export level1;
        interface fxp0;
    }
}
```

Trace IS-IS Protocol Traffic

To trace IS-IS protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] hierarchy level, and you can specify IS-IS–specific options by including the traceoptions statement at the [edit protocols isis] hierarchy level (for routing instances, include the statement at the [edit routing-instances *routing-instance-name* protocols isis] hierarchy level):

```
[edit protocols isis]
traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
}
```

You can specify the following IS-IS–specific options in the IS-IS flag statement:

all—Trace everything.

csn—Trace complete sequence number PDU (CSNP) packets.

error—Trace errored packets.

general—Trace general events.

hello—Trace hello packets.

lsp—Trace link-state PDU (LSP) packets.

lsp-generation—Trace link-state PDU generation packets.

normal—Trace normal events.

packets—Trace all IS-IS protocol packets.

policy—Trace policy processing.

psn—Trace partial sequence number PDU (PSNP) packets.

route—Trace routing information.

spf—Trace shortest-path-first (SPF) calculations.

state—Trace state transitions.

task—Trace routing protocol task processing.

timer—Trace routing protocol timer processing.

You can optionally specify one or more of the following flag modifiers:

detail—Detailed trace information

receive—Packets being received

send—Packets being transmitted

For information about tracing and global tracing options, see “Trace Global Routing Protocol Operations” on page 84.

Examples: Trace IS-IS Protocol Traffic

A common configuration traces SPF calculations, LSP calculations, normal protocol operations, and errors in protocol operation:

```
[edit]
protocols {
    isis {
        traceoptions {
            file isis-log size 1m files 10;
            flag spf;
            flag lsp;
            flag error;
            flag normal;
        }
    }
}
```

Trace only unusual or abnormal operations to the file routing-log, and trace detailed information about all IS-IS packets to the file isis-log:

```
[edit]
routing-options {
    traceoptions {
        file routing-log;
    }
}
protocols {
    isis {
        traceoptions {
            file isis-log size 10k files 5;
            flag csn detail;
            flag hello detail;
            flag lsp detail;
            flag psn detail;
        }
    }
}
```

Perform detailed tracing of mesh-group flooding:

```
[edit]
protocols {
    isis {
        traceoptions {
            file isis-log;
            flag lsp detail;
        }
    }
}
```

IS-IS LSP packets that contain errors are discarded by default. To log these errors, specify the error tracing operation:

```
[edit]
protocols {
    isis {
        traceoptions {
            file isis-log;
            flag error;
        }
    }
}
```

IS-IS Configuration Example

Configure IS-IS Level 1 only:

```
[edit]
protocols {
    isis {
        traceoptions {
            file isis-log size 1m files 10;
            flag spf;
            flag lsp;
            flag error;
            flag normal;
        }
        interface so-5/0/0 {
            level 2 {
                disable;
            }
        }
    }
}
```

Configure IS-IS Multicast Extensions

Most multicast routing protocols perform a reverse path forwarding (RPF) check on the source of multicast data packets. If a packet comes in on the interface that is used to send data to the source, the packet is accepted and forwarded to one or more downstream interfaces. Otherwise, the packet is discarded and a notification is sent to the multicast routing protocol running on the interface.

In certain instances, the unicast routing table used for the RPF check is also the table used for forwarding unicast data packets. Thus, unicast and multicast routing are congruent. In other cases, where it is preferred that multicast routing be independent of unicast routing, the multicast routing protocols are configured to perform the RPF check using an alternate unicast routing table inet.2.

You can configure IS-IS to calculate an alternate multicast topology, in addition to the normal unicast topology, and add the corresponding routes to inet.2. The IS-IS interface metrics for the multicast topology can be configured independently of the unicast metrics. You can also selectively disable interfaces from participating in the multicast topology while continuing to participate in the regular unicast topology. This lets you exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths.

The following sample configuration shows multicast topology support:

```
[edit]
   protocols {
      isis {
         traceoptions {
            file isis size 5m world-readable
         }
         multicast-topology;                # Enable multicast extensions
         interface so-0/0/0.0 {
            level 1 {
               metric 15;
               multicast-metric 18;     # Set Level 1 multicast metric
            }
            level 2 {
               metric 20;
               multicast-metric 14      # Set Level 2 multicast metric
               }
            }
         interface so-1/0/0.0 {
            level 1 {
               metric 15;
               multicast-metric 12;     # Set Level 1 multicast metric
            }
            level 2 {
               metric 20;
               multicast-metric 23      # Set Level 2 multicast metric
               }
            }
         interface so-2/0/0.0 {
            no-multicast;
            level 1 metric 14;
            level 2 metric 23;
         }
         interface fxp0.0 {
            disable;
         }
      }
   }
```